Towards Defeating the *Crossfire* Attack using SDN

Dimitrios Gkounis
FORTH, Greece
dgkounis@ics.forth.gr

Vasileios Kotronis
ETH Zurich, Switzerland
vkotroni@tik.ee.ethz.ch

Xenofontas Dimitropoulos
FORTH, Greece
fontas@ics.forth.gr

arXiv:1412.2013v1 [cs.NI] 5 Dec 2014

ABSTRACT

In this work, we propose online traffic engineering as a novel approach to detect and mitigate an emerging class of stealthy Denial of Service (DoS) link-flooding attacks. Our approach exploits the Software Defined Networking (SDN) paradigm, which renders the management of network traffic more flexible through centralised flow-level control and monitoring. We implement a full prototype of our solution on an emulated SDN environment using OpenFlow to interface with the network devices. We further discuss useful insights gained from our preliminary experiments as well as a number of open research questions which constitute work in progress.

1. INTRODUCTION

Some of the most powerful DoS attacks ever, such as the attack against Spamhaus which reached 400 Gbit/s [3], have been observed during the last two years. Moreover, new types of DoS link-flooding attacks have been recently proposed; these attacks are extremely difficult to mitigate and could potentially take entire countries off the Internet [8, 11]. The *Crossfire* [8] attack in particular: *i*) uses legitimate rather than spoofed source IP addresses to send traffic; these addresses are much harder to filter, *ii*) sends legitimate packets to publicly accessible servers, and *iii*) transmits low-bandwidth flows from each bot individually; these "under-the-radar" flows then cumulatively flood certain links in the network.

In this work, we introduce a novel approach to detect *Crossfire* bots under adverse DoS conditions. The key idea is to repeatedly change how traffic is routed and to monitor sources that persistently react to rerouting for the purpose of attacking a specific target. If certain sources are recorded several times in links that are DoS'ed they are considered suspicious and are filtered. Therefore, our approach effectively increases the cost for executing *Crossfire*. Traditionally, traffic engineering is used solely for the purpose of load-balancing DoS attacks. We use online traffic engineering both to detect stealthy bots and to balance the load. Our approach benefits from the following SDN principles: *i*) centralised visibility and control for fast traffic engineering, and *ii*) flow-level reactive traffic management.

2. METHODOLOGY

Attacker Model: The goal of the *Crossfire* [8] attack is to cut-off Internet connectivity towards a specific geographic area, the *Target Area*. Around the target area, there are public servers, the *Decoy Servers*, and network links, the *Target Links*, which lead to both the target area and the decoy servers. The attacker first constructs a map of the links (the *link-map*) in and around the target area. Then, he floods certain target links sending traffic only to the decoy servers. This way the flood cannot be directly observed in the target area, but it still isolates this area from the rest of the network. For this purpose, the attacker finds the target links that are used most along the bot-to-decoy server and bot-to-target area paths and he coordinates his bots to attack them. Finally, the attacker monitors the network routes and reacts to changes (which are possible defender's actions) by setting up the attack again, i.e., updating the link-map and recalculating the target links.

Defender Model: The goal of the defender is to keep the network running without any severe congestion and to find and rate-limit malicious traffic sources. Therefore, the defender: *i*) monitors traffic load on his network and detects DoS'ed links that are severely congested, *ii*) balances traffic load by rerouting traffic destined to different destinations (including the target area, decoy servers, etc), without knowing the attacker's classification, and *iii*) records sources observed in DoS'ed links to detect suspicious recurring sources.

Attacker - Defender Interplay: Malicious *Crossfire* bots will change their decoy server selection in case the rerouting has diverted their load from the targeted link(s), as shown in the example of Figure 1. Therefore, the same bots will "return to the crime-scene" and be present in another DoS event. The defender records the sources that appear in DoS'ed links. In addition, he records sources that he re-routes "away" of such links for load balancing. Sources that are pushed away, but return to future DoS'ed links (by selecting a new decoy server) are particularly suspicious. Thus after each attacker-defender interaction, malicious sources become more identifiable. Sources that surpass a threshold of

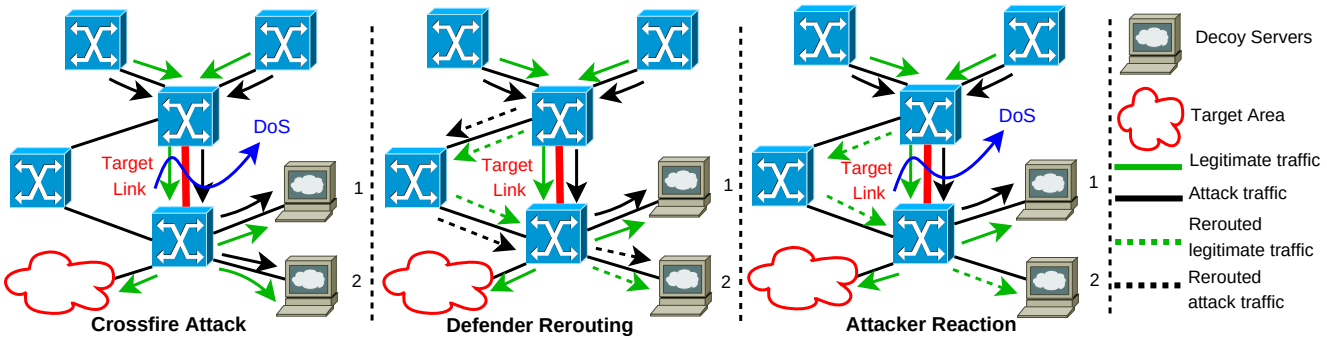

Figure 1: Attacker - Defender interplay: 1) Attacker floods target link, 2) Defender reroutes traffic destined to decoy server 2 to mitigate the attack, 3) Attacker updates selected decoy servers sending more traffic to decoy server 1. Traffic sources returning to the target link are deemed suspicious.

suspiciousness can be rate-limited. In contrast, benign flash-crowd load does not re-adjust to routing changes, but it uses the same popular destination(s) as before.

A key challenge for the defender is that he has no information about how the network nodes are mapped to the target area and the decoy servers. Therefore, he can select different re-routing strategies to optimise the detection phase: e.g., using random destinations or preferring servers that occupy homogeneous bandwidth on the congested links. The former strategy is simpler but might need more rounds to lead to detection. The latter strategy is a crude way of "approximating" the decoy server groups, since the attacker generally distributes traffic to those evenly for load-balancing.

Parameter Space: In our model, the attacker and the defender are continuously engaged into two co-dependent loops consisting of the following steps: 1) centralised flow monitoring [4]; 2) consolidation of measurement data and decision for next action; 3) network control and bot orchestration. The interactions between the two players determine if and when one of them has an advantage. Numerous parameters can influence the outcome, namely: *i*) network topology and link capacities, *ii*) attack flow rate per bot, *iii*) bot and flow assignment strategy for attacker, *iv*) link-map and link-bandwidth measurement intervals for both the attacker and the defender, *v*) rerouting strategy for the defender, *vi*) DoS'ed link detection approach and thresholds, and *vii*) suspiciousness threshold for rate-limiting decisions.

Proof-of-Concept Implementation: To better understand the parameter space and the parameterization trade-offs, we have implemented a full prototype of the attacker and the defender on the Mininet emulation environment [9]. Emulated bots periodically use traceroute towards the target area and the decoy servers to build and maintain the link-map. They flood the target links using low-bandwidth HTTP GET requests to the decoy servers. On the defender's side, we use the POX OpenFlow controller [2] for prototyping our online traffic engineering approach. We also employ OpenFlow [1] capabilities for network monitoring, in order to have a unified management solution. More implementation details of our emulation setup can be found in [7].

Preliminary Results and Insights: We evaluate the feasibility of our approach using a custom topology as described in [7]. Early results show that our SDN-based solution reacts in comparable time scales (some seconds) as the attack setup stage. Since the attacker needs approx. equal time to identify and react to defender's re-routing and counter-measures, the attack is effective for ∼50% of the total time. Moreover, a key insight is that links that are DoS'ed in parallel need to be handled in batches to avoid routing oscillations during mitigation; this mandates a tunable delay in the defender's action loop. We further note that the OpenFlow-based *control* functions run in sub-second times; *monitoring* is the dominant time component.

3. SUMMARY AND OPEN QUESTIONS

In this work, we proposed using dynamic traffic engineering in a novel way to detect and counter the particularly stealthy *Crossfire* attack. In contrast to the CoDef [10] approach which assumes a collaborative environment, we assume a hostile setup. The approach raises a number of research questions: *i*) What is the trade-off between detection accuracy, topology and traffic characteristics, and re-routing strategy? How can we re-route traffic to accelerate detection and minimise false positives? *ii*) What is the trade-off between rerouting costs (including possible instabilities) and DoS costs? *iii*) What are the temporal aspects of the attacker and defender control loops and their dependencies on different parameters? *iv*) How to extend the mitigation methodology to the inter-domain level involving support from upstream providers, who run the algorithm "as-a-Service"? *v*) How to scale up source recording using bloom filters and/or prefix aggregation? Other future research topics include understanding the interplay between fast rerouting and TCP congestion control [6], as well as the interaction with classic routing policies [5]. Lastly, we are interested in the game-theoretic modelling of the interaction between the players.

4. REFERENCES

[1] OpenFlow Switch Specification. `http://archive.openflow.org/documents/openflow-spec-v1.0.0.pdf`.

[2] POX. `http://www.noxrepo.org/pox/about-pox/`.

[3] The DDoS That Almost Broke The Internet. `http://blog.cloudflare.com/the-ddos-that-almost-broke-the-internet`.

[4] Athreya, A. P., Wang, X., Kim, Y. S., Tian, Y., and Tague, P. Resistance Is Not Futile: Detecting DDoS Attacks without Packet Inspection. In *Information Security Applications*. Springer, 2014, pp. 174–188.

[5] Caesar, M., and Rexford, J. BGP routing policies in ISP networks. *Network, IEEE 19*, 6 (2005), 5–11.

[6] Gao, R., Blair, D., Dovrolis, C., Morrow, M., and Zegura, E. Interactions of intelligent route control with TCP congestion control. In *NETWORKING 2007. Ad Hoc and Sensor Networks, Wireless Networks, Next Generation Internet*. Springer, 2007, pp. 1014–1025.

[7] Gkounis, D. Cross-domain DoS link-flooding attack detection and mitigation using SDN principles. Master's thesis, ETH Zurich, 2014.

[8] Kang, M. S., Lee, S. B., and Gligor, V. D. The crossfire attack. In *Security and Privacy (SP), 2013 IEEE Symposium on* (2013), IEEE, pp. 127–141.

[9] Lantz, B., Heller, B., and McKeown, N. A network in a laptop: rapid prototyping for software-defined networks. In *Proceedings of the 9th ACM SIGCOMM Workshop on Hot Topics in Networks* (2010), ACM, p. 19.

[10] Lee, S. B., Kang, M. S., and Gligor, V. D. CoDef: Collaborative Defense Against Large-scale Link-flooding Attacks. In *Proceedings of the Ninth ACM Conference on Emerging Networking Experiments and Technologies* (2013), CoNEXT '13.

[11] Studer, A., and Perrig, A. The Coremelt Attack. In *Proceedings of the 14th European Conference on Research in Computer Security* (2009), ESORICS'09.